

TSP with Time Windows and Service Time*

Yossi Azar[†]

Adi Vardi[‡]

January 27, 2015

Abstract

We consider TSP with time windows and service time. In this problem we receive a sequence of requests for a service at nodes in a metric space and a time window for each request. The goal of the online algorithm is to maximize the number of requests served during their time window. The time to traverse an edge is the distance between the incident nodes of that edge. Serving a request requires unit time. We characterize the competitive ratio for each metric space separately. The competitive ratio depends on the relation between the minimum laxity (the minimum length of a time window) and the diameter of the metric space. Specifically, there is a constant competitive algorithm depending whether the laxity is larger or smaller than the diameter. In addition, we characterize the rate of convergence of the competitive ratio to 1 as the laxity increases. Specifically, we provide a matching lower and upper bounds depending on the ratio between the laxity and the TSP of the metric space (the minimum distance to traverse all nodes). An application of our result improves the lower bound for colored packets with transition cost and matches the upper bound. In proving our lower bounds we use an interesting non-standard embedding with some special properties. This embedding may be interesting by its own.

*Supported in part by the Israel Science Foundation (grant No. 1404/10), by the Google Inter-university center and by The Israeli Centers of Research Excellence (I-CORE) program, (Center No.4/11).

[†]School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. E-mail: azar@tau.ac.il

[‡]School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. E-mail: adi.vardi@gmail.com

1 Introduction

Consider an employee in Google IT division. He is responsible for replacing malfunctioning disks in Google huge computer farms. During his shift he receives requests to replace disks at some points in time. Each request is associated with a deadline. If the disk will not be replaced before the deadline expiration, there is a high probability to a significant hit in the performance of the Search Engine. Replacing a disk takes a constant time (service time). However, before the employee can replace it, he must travel from his current location to the location of the disk. The goal is to maximize the number of disks replaced before their deadline expired. What path should the employee take and how the path changes with new requests? We call this problem TSP with time windows and service time (or vehicle routing with time windows and production costs). All the requests are of unit service time, and the window of request i is $[r_i, d_i]$. In this paper we characterize the competitive ratio for each metric space separately. We determine whether the competitive ratio is constant or not depending on the minimum laxity (the minimum length of a time window) and the diameter of the metric space (the maximum distance between nodes in the metric space). In addition, we consider the case where the laxity is large. Specifically, we provide a matching lower and upper bounds depending on the ratio between the laxity and the TSP of the metric space (the minimum distance to traverse all nodes).

Note that if the service time is negligible compared to the minimum positive distance between nodes then the problem reduces to the TSP (or vehicle routing) with time windows [4]. Moreover, if in addition all the deadlines are the same and all the release times are zero then the problem reduces to the well known (offline) orienteering problem [1, 2, 12].

We note that even when the service time is not negligible, the TSP with time windows and service time can be reduced to TSP with time windows [4] by changing the metric space. However, our competitive ratio depends on the properties of the metric space. The reduction might change the parameters of the metric space significantly. Hence, it might influence the crucial parameter which determines the competitive ratio. Therefore, we maintain the service time in our model.

Vehicle routing problems (with time windows) have been extensively studied both in computer science and operation research literature, see [11, 10, 15, 17, 16, 14]. For an arbitrary metric space Bansal et al. [4] showed $O(\log^2 n)$ -approximation (for certain cases a better approximation can be achieved [7]). Constant factor approximations have been presented for the case of points on a line [5, 18, 13]. For the orienteering problem, i.e., all the release times are zero and all the deadlines are the same, there are constant factor approximation algorithms [4, 8, 9, 6].

To the best of our knowledge we are the first to consider the online version of this general problem.

Another motivation for our problem is the colored Packets with Deadlines and Metric Space Transition Cost. In this problem we are given a sequence of incoming colored packets. Each colored packet is of unit size and has a deadline. There is a reconfiguration cost (setup cost) to switch between colors (the cost depends on the colors). The goal is to find a schedule that maximizes the number of packets that were transmitted before the deadline. Note that for one color the earliest deadline first (EDF) strategy is known to achieve an optimal throughput. The unit cost color has been considered in [3]. In particular, an application of our result to the uniform metric space, improves their lower bound and matches their upper bound.

1.1 Our results

Let $L = \min_{i \in \sigma} \{d_i - r_i\} \geq 1$ denote the minimum laxity of the requests (the minimum length of a time window). Let $\Delta(G)$ be the diameter of the metric space G , i.e., the largest distance between two nodes. Let $\text{TSP}(G)$ denote the weight of the minimal Traveling Salesperson Problem (TSP) in the metric space G . Let $\text{MST}(G)$ denote the weight of the minimal spanning tree (MST) in the metric space G .

In this paper we characterize when it is possible to achieve a $\Theta(1)$ competitive algorithm and when

the best competitive algorithm is unbounded. Moreover, we characterize the rate of convergence of the competitive ratio to 1 as the laxity increases. Specifically, we provide a matching lower and upper bounds depending on the ratio between the laxity and the TSP of the metric space. We consider three cases.

- Case A: $L < \Delta(G)/2$. For any metric space the competitive ratio of any online algorithm is unbounded. The claim is easily proved.
- Case B: $(2 + \epsilon)\Delta(G) < L \leq TSP(G)$ for any $\epsilon > 0$. We design $O(1)$ competitive algorithm and a 1.002 lower bound.
- Case C: $L > TSP(G)$. Let $\delta = TSP(G)/L < 1$. We show a strictly larger than 1 lower bound. Specifically, if $\delta \leq \frac{1}{9}$ we provide a lower bound of $1 + \Omega(\sqrt{\delta})$ as well as a matching upper bound of $1 + O(\sqrt{\delta})$. We note that without service time it is easy to design 1-competitive algorithm by traveling over TSP periodically. Recall that there is a reduction from the service time model to a model without service time that seems to contradict the lower bound (see [4]). Nevertheless, the reduction modifies the metric space and hence increases δ such that δ is not smaller than $\frac{1}{9}$.

Note that in the remaining cases, i.e., $\Delta(G)/2 \leq L \leq (2 + \epsilon)\Delta(G)$ the question whether there exists a constant competitive algorithm depends on the metric space. Specifically, for $L = \Delta(G)$ it is easy to prove that there is no constant competitive algorithm for the uniform metric space. In contrast, there is a constant competitive algorithm for the line metric space.

Observe that when the metric space consists of a single node (i.e., no traveling time) the optimal algorithm to serve a requests is EDF (earliest deadline first) which is 1-competitive. This case is equivalent to packet scheduling with deadlines. If packets have colors and switching between colors costs 1, then our result improves the lower bound of [3]. Specifically, we improve their $1 + \Omega(\delta)$ lower bound to $1 + \Omega(\sqrt{\delta})$ (where C is the number of nodes in G and hence, $\delta = \frac{C}{L} < 1$) and match their upper bound for the uniform metric space.

It is also interesting to mention that in many cases the competitive ratio of an algorithm is computed as a supremum over all the metric spaces and the lower bound is proved for one specific metric space. We prove more refined results. Specifically, we show an upper bound and a lower bound for each metric space separately. Hence, one can not design a better competitive algorithm for the specific metric space that one encounter in the real specific instance.

Embedding result. One of the technique that is used for the lower bound is an embedding that is interesting on its own. Let $w(S)$ denote the weight of the star metric S (i.e., the sum of the weights of the edges of S). We prove that for any given metric space G on nodes V and for any vertex $v_0 \in V$ there exists a star metric S with leaves V and an embedding $f : G \rightarrow S$ from G to S (f depends on v_0) such that:

1. $w(S) = MST(G)$.
2. The weight of every Steiner tree in S that contains v_0 is not larger than the weight of the Steiner tree on the same nodes in G .

Note that this embedding is different from the usual embedding since we do not refer specifically to distances between vertices. Typically, embedding is used to prove an upper bound by simplifying the metric space. In contrast, our embedding is used to prove a lower bound.

In order to prove the lower bound we first establish it for a star metric, and then extend it for a general metric space. Note that a lower bound on a sub-graph is not a lower bound on the ambient graph. For example, a lower bound for MST of a metric space G is not a lower bound for G since the algorithm may use the additional edges to reduce the transition time.

2 The Model

We formally model the TSP with time windows and service time problem as follows. Let $G = (V, w)$ be a given metric space where V is a set of n nodes and w is a distance function. We are given an online sequence of requests for service. Each request is characterized by a pair $([r_i, d_i], v_i)$, where $r_i \in N_+$ and $d_i \in N_+$ are the respective arrival time and deadline time of the request, and $v_i \in V$ is a node in the metric space G . The time to traverse from node v_i to node v_j is $w(v_i, v_j)$. Serving a request at some node requires unit size service time. The goal is to serve as many requests as possible within their time windows $[r_i, d_i]$.

When all r_i are equal to 0 and all d_i are equal to B and the service time is negligible the problem is reduced to the well-known orienteering problem with budget B and a prize for each node which is equal to the number of requests in this node. That is, finding a path of total distance at most B that maximizes the prize of all visited nodes.

Let $\text{ALG}(\sigma)$ ($\text{OPT}(\sigma)$) denote the throughput of the online (respectively, optimal) algorithm with respect to a sequence σ . We consider the benefit problem and hence $\inf_{\sigma} \text{OPT}(\sigma) / \text{ALG}(\sigma) \geq 1$.

3 Embedding of Metric Spaces

In this section we describe an embedding of a general metric space into a star metric. We begin by introducing some new definitions:

- We define $w(T) = \sum_{e \in E} w(e)$ for a tree $T = (V, E)$, and let $P_T(v)$ denote the parent of node v in a rooted tree T .
- Let S be a star metric with a center c . Let w_i denote the weight of the edge incident to the vertex v_i . We define $w_S(V) = \sum_{v \in V} w(c, v) (= \sum_{v_i \in V} w_i)$. It is clear that for a star S with leaves V , $w_S(V) = w(S)$.
- Let $T_G(V)$ be the minimum weight connected component that contains the set V (i.e., the minimum Steiner tree on these points) in the metric space G .

Note that $E(G)$ is the set of edges of graph G . Recall that $\text{MST}(G)$ denote the weight of the minimal spanning tree (MST) in the metric space G .

Theorem 3.1 *For any given metric space G on nodes V and for any vertex $v_0 \in V$ there exists a star metric S with leaves V and an embedding $f : G \rightarrow S$ from G to S (f depends on v_0) such that:*

1. $w(S) = w(T_G(V)) (= \text{MST}(G))$.
2. For every $V' \subseteq V$ such that $v_0 \in V'$, $w(T_G(V')) \geq w_S(V')$.

Proof. We prove the theorem by describing a star metric with the required properties. Let G be a given metric space on nodes V and a leaf $v_0 \in V$. Let T be the MST for G created by means of Prim's algorithm with the root v_0 . Let S be a star metric with leaves V such that for each $u \in V$, $w_u = w(u, P_T(u))$. Clearly, $w_{v_0} = 0$. We prove that S and v_0 satisfy the theorem's properties:

Property 1: Clearly, $w(S) = w(T)$, and since T is a MST for G , $w(S) = w(T) = \text{MST}(G) = w(T_G(V))$.

Property 2: We have to prove that for every $V' \subseteq V$ such that $v_0 \in V'$, $w(T_G(V')) \geq w_S(V')$. Let $V' = \{v_0, v_{i_1}, \dots, v_{i_{r-1}}\}$. Recall that we defined $w_u = w(u, P_T(u))$. Clearly, $w_S(V') = \sum_{j=1}^{r-1} w(v_{i_j}, P_T(v_{i_j}))$. Hence it suffices to prove that $w(T_G(V')) \geq \sum_{j=1}^{r-1} w(v_{i_j}, P_T(v_{i_j}))$. The proof is based on the following idea. We begin with the minimum Steiner tree that contains V' (meaning $T_G(V')$). Then we transform it to an MST on all vertices by running Prim from v_0 and replacing the Steiner tree's edges with Prim's edges.

We prove that each time the algorithm adds an edge e that corresponds to an edge in $w_S(V')$ it deletes an edge e' from $T_G(V')$ such that $w(e) \leq w(e')$. Note that we also add edges incident to vertices not in V' in order to maintain a tree. The weights of these edges are ignored. Since the algorithm starts with $T_G(V')$ and finishes with T , this proves that the property holds (recall that the weight of the edges of S is determined by the weight of the edges of T). The exact description of our algorithm, called the Embed-Prim algorithm, is provided in Figure 1.

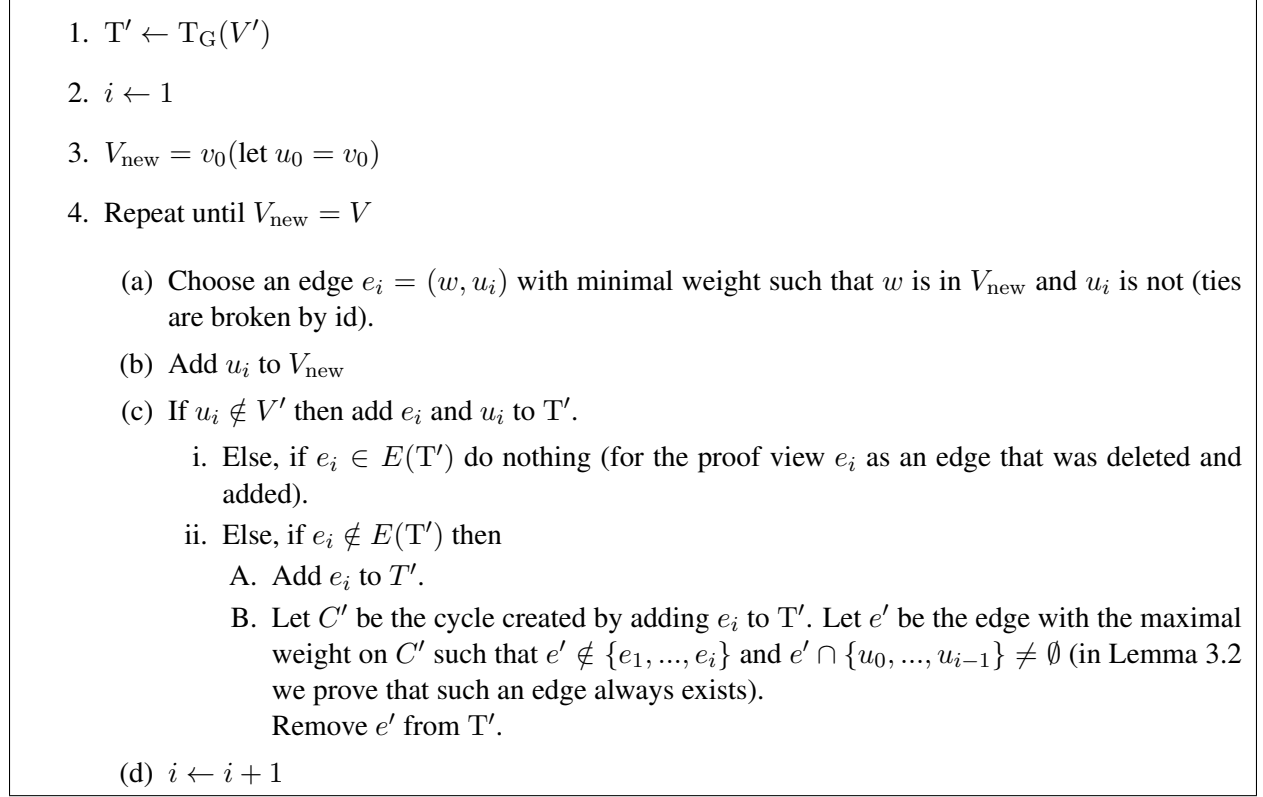


Figure 1: Algorithm Embed-Prim.

First we show the correctness of Embed-Prim:

Lemma 3.2 *Let C' be the cycle created in step 4(c)ii in the algorithm. There exists at least one edge e' that belongs to C' , such that $e' \notin \{e_1, \dots, e_i\}$ and $e' \cap \{u_0, \dots, u_{i-1}\} \neq \emptyset$.*

Proof. Note that a cycle is created in step 4(c)ii in the algorithm since adding an edge to a tree always creates a cycle. Similar to Prim, the edges that Embed-Prim adds after step 1 in the algorithm do not create a cycle. Therefore C' must contain edges added in step 1 in the algorithm. At least one of these edges must touch one of the vertices $\{u_0, \dots, u_{i-1}\}$ \square

Lemma 3.3 *After each step of Embed-Prim, T' is a tree which contains V' .*

Proof. At the beginning T' is $T_G(V')$, which is a tree that contains V' . We never remove vertices and hence T' always contains V' . Whenever we add an edge that creates a cycle we open the cycle by removing an edge from it. \square

Now we claim that Embed-Prim satisfies the following invariant:

Lemma 3.4 *Each time Embed-Prim adds an edge e that corresponds to an edge in $w_S(V')$, it deletes an edge e' from $T_G(V')$ such that $w(e) \leq w(e')$.*

Proof. Step 4c in the algorithm is irrelevant, since the edge does not correspond to an edge in $w_S(V')$ (the vertex that was added by Embed-Prim is not in V'). In step 4(c)i in the algorithm, $w(e) = w(e')$. In step 4(c)ii in the algorithm, since Embed-Prim could have added edge e' , but did choose the edge e instead, $w(e) \leq w(e')$ (recall that Embed-Prim always chooses the edge with the minimal weight). \square

Now we are ready to prove that S satisfies the second property of the embedding. By the definition of Prim $e_i = (u_i, P_T(u_i))$. Hence, $\sum_{j=1}^{r-1} w(e_{i_j}) = \sum_{j=1}^{r-1} w(v_{i_j}, P_T(v_{i_j}))$. Let e'_i be the edge deleted from T' when edge e_i was added (steps 4(c)i, 4(c)ii in the algorithm). Then

$$w_S(V') = \sum_{j=1}^{r-1} w(v_{i_j}, P_T(v_{i_j})) = \sum_{j=1}^{r-1} w(e_{i_j}) \leq \sum_{j=1}^{r-1} w(e'_{i_j}) \leq w(T_G(V')).$$

where the first equality follows from the definition, the first inequality results from the invariance, and the last inequality follows from the definition. \square

4 Lower Bounds

4.1 Lower Bound for a Large Diameter Laxity Ratio (Case A)

In this section we consider the case where $L < \Delta(G)/2$ (recall that $\Delta(G)$ is the diameter and L is the laxity) then we show that the competitive ratio of any algorithm is unbounded.

Theorem 4.1 *No online algorithm can achieve a bounded competitive ratio for any metric space in which $L < \Delta(G)/2$.*

Proof. Let G be any metric space. Every $\Delta(G) + 1$ units of time we bring a request with laxity of L on a node which is at a distance of at least $\Delta(G)/2$ from the current location of the online algorithm. It is clear that the algorithm can not serve any requests while OPT can serve all the requests. \square

4.2 Lower Bound for a Small Diameter Laxity Ratio (Case B and C)

In this section we consider Cases B and C. Let $\delta = TSP(G)/L$. If $\delta < 1$ (Case C) We show a strictly larger than 1 lower bound. Specifically, if $\delta \leq \frac{1}{9}$ we provide a lower bound of $1 + \Omega(\sqrt{\delta})$. If $\delta > \frac{1}{9}$ (Case B) we can use requests with laxity of $9TSP(G)$ (i.e., $\delta = \frac{1}{9}$), and obtain a lower bound of 1.002. Therefore, from now on we only consider Case C.

4.2.1 Lower Bound for a Star Metric

In this section we consider the case where the traveling time between nodes is represented by a star metric. This is also equivalent to the case where the traveling time from node i is w_i .

The general idea is that the adversary creates many requests with large deadline at node v_0 at each time unit, and also blocks of fewer requests with close deadlines at other nodes. Any online algorithm must choose between serving many requests with large deadline or traveling between many of the nodes and serving also the requests with close deadline.

Recall that $w(S)$ denote the weight of the star metric S (i.e., the sum of the weights of the edges of S). We define $F = \sqrt{w(S)L}$. Let $\delta = TSP(G)/L = w(S)/L$.

Theorem 4.2 *No deterministic or randomized online algorithm can achieve a competitive ratio better than $1 + \Omega(\sqrt{\delta})$ in any given star metric S for $\delta \leq \frac{1}{9}$. Otherwise, if $\delta > \frac{1}{9}$, the bound becomes 1.002.*

Proof. Let S be a given star metric with nodes $V = v_0, \dots, v_{n-1}$. We can assume, without loss of generality, that $\delta \leq \frac{1}{9}$, since otherwise one may use requests with laxity of $9w(S)$ (i.e., $\delta = \frac{1}{9}$), and obtain a lower bound of 1.002. Let type A node denote node $v_0 \in V$ and type B node denote nodes $v_1, \dots, v_{n-1} \in V$. Let type A request and type B request refer to requests with type A node and type B node, respectively. Recall that w_i denote the weight of the edge incident to the vertex v_i .

We begin by describing the sequence $\sigma(S, \text{ALG})$.

Sequence structure: Recall that each request is characterized by a pair $([r_i, d_i], v_i)$, where $r_i \in N_+$ and $d_i \in N_+$ are the respective arrival time and deadline time of the request, and v_i is a node in S . There are up to $N = \frac{L}{3F} = \frac{1}{3} \sqrt{\frac{L}{w(S)}}$ blocks, where each block consists of $3F$ time units. Let $t_i = 1 + 3(i-1)F$ denote the beginning time of block i . For each block i , where $1 \leq i \leq N$, F requests located at various nodes arrive at the beginning of the block. Specifically, $\frac{w_j}{w(S)-w_0} F$ type B requests $([t_i, L+t_i], v_j)$, for each $1 \leq j \leq n-1$ are released. A type A request $([t, 3L], v_0)$ is released at each time unit t in each block. Once the adversary stops the blocks, additional requests arrive (we call this the final event). The exact sequence is defined as follows:

1. $i \leftarrow 1$
2. Add block i
3. If with probability at least $1/4$ there are at least $F/2$ unserved type B requests at the end of block i (denoted by Condition 1), then L requests $([t_{i+1}, L+t_{i+1}], v_1)$ are released and the sequence is terminated. See Figure 5. Clearly, t_{i+1} is the time of the final event. Denote this by Termination Case 1.
4. Else, if with probability at least $1/4$, at most $2F$ requests were served during block i (denoted by Condition 2), then $3L$ requests $([t_{i+1}, 3L], v_0)$ are released and the sequence is terminated. Clearly, t_{i+1} is the time of the final event. See Figure 6. Denote this by Termination Case 2.
5. Else, if $i = N$ (there are N blocks, none of them satisfied Condition 1 or 2) $3L$ requests $([L+1, 3L], v_0)$ are released, and the sequence is terminated. Clearly, $L+1$ is the time of the final event. See Figure 7. Denote this by Termination Case 3.
6. Else ($i < N$) then $i \leftarrow i+1$, Goto 2

We make the following **observations**: (i) Each block consists of $3F$ time units. Hence, if ALG served at most $2F$ requests during a block, there must have been at least F idle time units. (ii) There are up to $\frac{1}{3} \sqrt{\frac{L}{w(S)}}$ blocks and each block consists of $3\sqrt{w(S)L}$ time units. Hence, the time of the final event is at most $L+1$. (iii) Exactly one type A request arrives at each time-slot until the final event. Hence, at most L type A requests arrive before (not including) the final event. (iv) During each block, exactly F type B requests arrive, which sum up to at most $L/3$ type B requests before (not including) the final event.

Now we can analyze the competitive ratio of $\sigma(S, \text{ALG})$. Consider the following possible sequences (according to the termination type):

1. Termination Case 1: Let Y denote the number of requests in the sequence. According to the observations, the sequence consists of at most L type A requests, and at most $\frac{4}{3}L$ type B requests ($L/3$ until the final event and L at the final event). Hence, $Y \leq L + \frac{4}{3}L \leq 3L$.

- **We bound the performance of ALG:** At time t_{i+1} there is a probability of at least $1/4$ that ALG has $L + F/2$ unserved type B requests. Since type B requests have laxity of L , ALG can serve at most $L + 1$ of them, and must drop at least $F/2 - 1$. The expected number of served requests is

$$E(\text{ALG}(\sigma)) \leq Y - \frac{1}{4}(F/2 - 1) = Y - \frac{1}{8}F + 1/4.$$

- **We bound the performance of OPT':** OPT' serves the requests in three stages:
 - **Type B requests that arrive before the final event:** Recall that all type B requests in a block arrive at once in the beginning of the block. In each block OPT' serves first all the requests to node v_1 , then all the requests to node v_2 , and so on. It is clear that OPT' needs at most $F + 2w(S)$ time units to serve the requests (F for serving and $2w(S)$ for traveling). OPT' serves the requests starting from the beginning of the block. Recall that $L \geq 9w(S)$ and $F = \sqrt{w(S)L}$. Therefore $2F \geq 18w(S)$. Since the block's size is $3F$, there are enough time units. Moreover, since $L \geq 9w(S)$, $L \geq 3\sqrt{w(S)L} = 3F > F + 2w(S)$. Hence, all the requests can be served before deadline expiration.
 - **Type B requests that arrive during the final event:** The L requests $([t_{i+1}, L + t_{i+1}], v_1)$ arrived during the final release time are served by OPT' consecutively from time t_{i+1} . OPT' can serve L requests, except for one travel phase, and hence may lose at most $2w(S)$ requests. According to the observations, the time of the final event t_{i+1} is at most $L + 1$. We conclude that OPT' serves all type B requests until time unit $2L$.
 - **Type A requests:** OPT' serves the L type A requests consecutively from time unit $2L + 1$. Since the deadlines are $3L$, OPT' serves all type A requests.

We conclude that $\text{OPT}(\sigma) \geq \text{OPT}'(\sigma) \geq Y - 2w(S)$.

The competitive ratio is

$$\frac{\text{OPT}(\sigma)}{E(\text{ALG}(\sigma))} \geq \frac{Y - 2w(S)}{Y - \frac{1}{8}F + 1/4} \geq \frac{3L - 2w(S)}{3L - \frac{1}{8}(\sqrt{w(S)L}) + 1/4} = 1 + \Omega\left(\sqrt{\frac{w(S)}{L}}\right).$$

Here the second inequality results from the fact that the number is above 1 and the numerator and the denominator increase by the same value.

2. Termination Case 2: The sequence consists of more than $3L$ type A requests, all deadlines are at most $3L$.

- **We bound the performance of ALG:** The probability that ALG was idle during F time units is at least $1/4$. Hence, the expected number of served requests is $E(\text{ALG}(\sigma)) \leq 3L - \frac{1}{4}F$.
- **We bound the performance of OPT':** At each time unit until the final event, OPT' serves the type A request that arrived at that particular time unit. Consequently, from the final event and until time unit $3L$, OPT' serves the type A requests that arrived at the final event. Therefore, OPT' serves $3L$ type A requests, and so $\text{OPT}(\sigma) \geq \text{OPT}'(\sigma) \geq 3L$.

The competitive ratio is

$$\frac{\text{OPT}(\sigma)}{E(\text{ALG}(\sigma))} \geq \frac{3L}{3L - \frac{1}{4}F} = \frac{3L}{3L - \frac{1}{4}(\sqrt{w(S)L})} = 1 + \Omega\left(\sqrt{\frac{w(S)}{L}}\right).$$

3. Termination Case 3: the sequence consists of $3L$ type A requests, all deadlines are at most $3L$.

- **We bound the performance of ALG:** Let U_i be the event that the number of unserved type B requests at the end of block i is less than $F/2$. If U_i occurs, then let j_k , $1 \leq k \leq r$, be the type B nodes visited by ALG in block i . At least $F/2$ requests that arrived in this block have to be served (recall that F type B requests arrive at the beginning of each block). Therefore,

$$\frac{w_{j_1}}{w(S) - w_0}F + \frac{w_{j_2}}{w(S) - w_0}F + \cdots + \frac{w_{j_r}}{w(S) - w_0}F \geq F/2,$$

and so

$$w_{j_1} + w_{j_2} + \cdots + w_{j_r} \geq \frac{w(S) - w_0}{2}.$$

Let E_i be the event that more than $2F$ requests are served during block i . If event U_{i-1} and E_i occur, then there are at most $3F/2$ unserved type B requests in the beginning of block i (F arrived in the beginning of the block and at most $F/2$ from the previous block) but more than $2F$ requests were served. Therefore, at least one type A request was served during the block. Combining the results, if U_i , U_{i-1} and E_i occur then:

- During block i at least $(w(S) - w_0)/2$ time units were used for traveling between type B nodes.
- Type A request was served during the block.

A block i is called *good* if the events U_i , U_{i-1} and E_i occur. For any two (consecutive) good blocks the traveling cost is at least $(w(S) - w_0)/2 + w_0 \geq w(S)/2$. Since none of the blocks satisfy Condition 1 or 2, it follows that for all i such that $\frac{1}{3}\sqrt{\frac{L}{w(S)}} \geq i \geq 1$ we have: $\Pr[U_i] \geq 3/4$, $\Pr[U_{i-1}] \geq 3/4$, and $\Pr[E_i] \geq 3/4$. Therefore:

$$\begin{aligned} \Pr[U_i \cap U_{i-1} \cap E_i] &= 1 - \Pr[\neg(U_i \cap U_{i-1} \cap E_i)] \\ &= 1 - \Pr[\neg U_i \cup \neg U_{i-1} \cup \neg E_i] \geq 1 - 1/4 - 1/4 - 1/4 = 1/4. \end{aligned}$$

The sequence consists of $\frac{1}{3}\sqrt{\frac{L}{w(S)}}$ blocks. Therefore, the expected number of good blocks is $\frac{1}{4} \cdot \frac{1}{3}\sqrt{\frac{L}{w(S)}} = \frac{1}{12}\sqrt{\frac{L}{w(S)}}$ and hence the expected number of disjoint pairs of blocks is $\frac{1}{24}\sqrt{\frac{L}{w(S)}}$. Consequently, the expected number of lost requests is at least $\frac{1}{24}\sqrt{\frac{L}{w(S)}} \frac{w(S)}{2}$. We conclude that the expected number of served requests is

$$E(\text{ALG}(\sigma)) \leq 3L - \frac{1}{48}w(S)\sqrt{\frac{L}{w(S)}} = 3L - \frac{1}{48}\left(\sqrt{w(S)L}\right).$$

- **We bound the performance of OPT':** At each time unit until the final event, OPT' serves the type A request that arrived at the same time unit. Consequently, from the final event and until time unit $3L$, OPT' serves the type A requests that arrived at the final event. Therefore, OPT' serves $3L$ type A requests, and so $\text{OPT} \geq \text{OPT}' \geq 3L$.

The competitive ratio is

$$\frac{\text{OPT}(\sigma)}{E(\text{ALG}(\sigma))} \geq \frac{3L}{3L - \frac{1}{48}\left(\sqrt{w(S)L}\right)} = 1 + \Omega\left(\sqrt{\frac{w(S)}{L}}\right).$$

Note that in all 3 cases we got $1 + \Omega\left(\sqrt{\frac{w(S)}{L}}\right) = 1 + \Omega(\sqrt{\delta})$. This completes the proof.

□

Corollary 4.3 *No deterministic or randomized online algorithm can achieve a competitive ratio better than $1 + \Omega(\sqrt{n/L})$ when all traveling times takes one unit of time and $L \geq 9n$. Otherwise, if $L < 9n$, the bound becomes 1.002 (note that in this case $\delta = n/L$).*

Proof. Let S be a star metric with n nodes such that the weight of each edge is equal to $1/2$. Clearly, traveling between each two nodes requires one time unit and $w(S) = n/2$. Applying Theorem 4.2, we obtain the lower bound of $1 + \Omega(\sqrt{n/L})$. □

Observe that if $n = C$ the bound becomes $1 + \Omega(\sqrt{C/L})$ which improves the lower bound of $1 + \Omega(C/L)$ from [3] (since $C < L$).

4.2.2 Lower Bound for a General Metric Space

In this section we consider the case where the traveling time between nodes is represented by a metric space G . Note that a lower bound for a star metric space does not imply a lower bound for a general metric space. Recall that $\delta = \text{TSP}(G)/L < 1$.

We use the embedding from Theorem 3.1 to prove a $1 + \Omega(\sqrt{\delta})$ lower bound.

Theorem 4.4 *No deterministic or randomized online algorithm can achieve a competitive ratio better than $1 + \Omega(\sqrt{\delta})$ in any given metric space G , for $\delta \leq \frac{1}{9}$. Otherwise, if $\delta > \frac{1}{9}$, the bound becomes 1.002.*

5 Upper bounds

5.1 Constant Approximation Algorithm for case B

In this section we design a deterministic online algorithm, for a general metric space where $L > 9\Delta(G)$ (recall that $\Delta(G)$ is the diameter of G). The algorithm achieves a constant competitive ratio. As shown in the previous section, no online algorithm can achieves a competitive ratio better 1.002. A more precise analysis can replace $L > 9\Delta(G)$ with $L > (2 + \epsilon)\Delta(G)$ for any $\epsilon > 0$.

The algorithm which we call ORIENT-WINDOW combines the following ideas.

- The algorithm works in phase of $K = 3\Delta(G)$. In each phase the algorithm serves only requests that arrived in the previous phases, and will not expired during the phase. Due to this perturbation we lose a constant factor.
- The decision which requests will be served in a phase ignore their deadlines. Due to this violation of EDF we lose a constant factor.
- In each phase the algorithm serves requests node by node. The order of the nodes is determined by solving an orienteering problem. Since a constant approximation algorithm is known to the orienteering problem, we lose a constant factor.

Theorem 5.1 *The algorithm ORIENT-WINDOW attains a competitive ratio of $O(1)$.*

In each phase $\ell = 1, 2, \dots$, do

- Beginning of the phase (at time $K(\ell - 1)$)
 - Decrease the deadline of each unserved request (r, d, v) from d to $K \lfloor d/K \rfloor$.
 - Let R^ℓ be the collection of unserved requests such that their decreased deadline was not exceeded. Let $S_j^\ell \subseteq R^\ell$ denote the subset of requests at node v_j in R^ℓ .
 - Using a constant approximation algorithm solve the unrooted orienteering problem with budget $\Delta(G)$ where the prize of a node v_j is the number of requests in S_j^ℓ . Let $v_{i_1}, v_{i_2}, \dots, v_{i_r}$ denote the order of the nodes in the solution.
 - ρ_ℓ consists of all requests of $S_{i_1}^\ell$ scheduled consecutively, then all requests of $S_{i_2}^\ell$ scheduled consecutively, and so on.
- During the phase (between time $K(\ell - 1)$ and time $K\ell$)
 - The requests are served according to ρ_ℓ (unserved requests in the suffix of ρ_ℓ , due to the end of the phase are dropped).

Figure 2: Algorithm ORIENT-WINDOW

5.2 Asymptotically Optimal Algorithm for Case C

In this section we design a deterministic online algorithm, for a general metric space. The algorithm achieves a competitive ratio of $1 + o(1)$ when the minimum laxity of the requests is asymptotically larger than the weight of the TSP (as shown in the previous sections, this is essential).

The algorithm is a natural extension of the BG algorithm from [3]. Our algorithm, which we call TSP-EDF, formally described in Figure 3, works in phases of $K = \sqrt{\text{TSP}(G)L}$ time units. In each phase the algorithm serves requests node by node. The order of the nodes is determined by the minimum TSP or an approximation. The algorithm achieves a competitive ratio of $1 + O(\sqrt{\text{TSP}(G)/L})$ for $L > 10\text{TSP}(G)$.

In each phase $\ell = 1, 2, \dots$, do

- Beginning of the phase (at time $K(\ell - 1)$)
 - Decrease the deadline of each unserved request (r, d, v) from d to $K \lfloor d/K \rfloor$.
 - Let R^ℓ be the collection of unserved requests such that their decreased deadline was not exceeded. Let S^ℓ be the K -length prefix of EDF (earliest deadline first) schedule (according to the modified deadline) of R^ℓ . Let $S_j^\ell \subseteq S^\ell$ denote the subset of requests at node v_j in S^ℓ . Let $v_{i_1}, v_{i_2}, \dots, v_{i_n}$ denote the order of the nodes in the minimal TSP (or approximation).
 - ρ_ℓ consists of all requests of $S_{i_1}^\ell$ served consecutively, then all requests of $S_{i_2}^\ell$ served consecutively, and so on.
- During the phase (between time $K(\ell - 1)$ and time $K\ell$)
 - The requests are served according to ρ_ℓ (unserved requests in the suffix of ρ_ℓ , due to the end of the phase are dropped).

Figure 3: Algorithm TSP-EDF.

Theorem 5.2 *The algorithm TSP-EDF attains a competitive ratio of $1 + O(\sqrt{\text{TSP}(G)/L})$.*

References

- [1] E. M. Arkin, J. S. Mitchell, and G. Narasimhan. Resource-constrained geometric network optimization. In *Proceedings of the fourteenth annual symposium on Computational geometry*, pages 307–316. ACM, 1998.
- [2] B. Awerbuch, Y. Azar, A. Blum, and S. Vempala. New approximation guarantees for minimum-weight k-trees and prize-collecting salesmen. *SIAM Journal on Computing*, 28(1):254–262, 1998.
- [3] Y. Azar, U. Feige, I. Gamzu, T. Moscibroda, and P. Raghavendra. Buffer management for colored packets with deadlines. SPAA '09, pages 319–327. ACM, 2009.
- [4] N. Bansal, A. Blum, S. Chawla, and A. Meyerson. Approximation algorithms for deadline-tsp and vehicle routing with time-windows. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 166–174. ACM, 2004.
- [5] R. Bar-Yehuda, G. Even, and S. M. Shahar. On approximating a geometric prize-collecting traveling salesman problem with time windows. *Journal of Algorithms*, 55(1):76–92, 2005.
- [6] A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 46–55. IEEE, 2003.
- [7] C. Chekuri and N. Korula. Approximation algorithms for orienteering with time windows. *arXiv preprint arXiv:0711.4825*, 2007.
- [8] C. Chekuri, N. Korula, and M. Pál. Improved algorithms for orienteering and related problems. *ACM Transactions on Algorithms (TALG)*, 8(3):23, 2012.
- [9] C. Chekuri and A. Kumar. Maximum coverage problem with group budget constraints and applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 72–83. Springer, 2004.
- [10] M. Desrochers, J. Desrosiers, and M. Solomon. A new optimization algorithm for the vehicle routing problem with time windows. *Operations research*, 40(2):342–354, 1992.
- [11] M. Desrochers, J. K. Lenstra, M. W. Savelsbergh, and F. Soumis. Vehicle routing with time windows: Optimization and approximation. *Vehicle routing: Methods and studies*, 16:65–84, 1988.
- [12] B. L. Golden, L. Levy, and R. Vohra. The orienteering problem. *Naval research logistics*, 34(3):307–318, 1987.
- [13] Y. Karuno and H. Nagamochi. 2-approximation algorithms for the multi-vehicle scheduling problem on a path with release and handling times. *Discrete Applied Mathematics*, 129(2):433–447, 2003.
- [14] H. Nagamochi and T. Ohnishi. Approximating a vehicle scheduling problem with time windows and handling times. *Theoretical Computer Science*, 393(1):133–146, 2008.
- [15] M. W. Savelsbergh. Local search in routing problems with time windows. *Annals of Operations Research*, 4(1):285–305, 1985.
- [16] K. C. Tan, L. H. Lee, Q. Zhu, and K. Ou. Heuristic methods for vehicle routing problem with time windows. *Artificial intelligence in Engineering*, 15(3):281–295, 2001.

- [17] S. R. Thangiah. *Vehicle routing with time windows using genetic algorithms*. Citeseer, 1993.
- [18] J. N. Tsitsiklis. Special cases of traveling salesman and repairman problems with time windows. *Networks*, 22(3):263–282, 1992.

A Proofs

A.1 Proof of Theorem 4.4

Let G be a given metric space on nodes V . We use the embedding from Theorem 3.1. Let S, v_0 be the output of the embedding. Let $\sigma(S, \text{ALG})$ be the sequence described in Theorem 4.2, when v_0 is type A node and the other nodes are type B. Recall that, by definition, $F = \sqrt{w(S)L}$. We use σ for ALG on G . We can assume, without loss of generality, that $\delta \leq \frac{1}{9}$ since otherwise one may use requests with laxity of $9TSP(G)$ (i.e., $\delta = \frac{1}{9}$), and obtain a lower bound of 1.002. Note that for any metric space G we have $MST(G) \leq TSP(G) < 2MST(G)$. Consider the following possible cases, similar to the proof of Theorem 4.2.

1. In Termination Case 1 there exists a block i such that, with probability at least $1/4$, at the end of the block there are at least $F/2$ unserved type B requests. In Theorem 4.2 we proved that:
 - The sequence consists of up to $3L$ requests.
 - The expected number of requests ALG missed is at least $F/8 - 1/4$.
 - OPT missed up to $TSP(G)$ requests (while transmitting the type B requests that arrived during the final event).

Therefore, the competitive ratio depends only on F , $TSP(G)$ and L :

$$\begin{aligned} \frac{\text{OPT}(\sigma)}{E(\text{ALG}(\sigma))} &\geq \frac{3L - TSP(G)}{3L - \frac{1}{8}F + 1/4} = \frac{3L - TSP(G)}{3L - \frac{1}{8}(\sqrt{w(S)L}) + 1/4} \\ &= \frac{3L - TSP(G)}{3L - \frac{1}{8}(\sqrt{MST(G)L}) + 1/4} = 1 + \Omega\left(\sqrt{TSP(G)/L}\right). \end{aligned}$$

Here the second equality results from the fact that $w(S) = MST(G)$.

2. In Termination Case 2 there exists a block i such that, with probability at least $1/4$, at most $2F$ requests were served during the block. In Theorem 4.2 we proved that:
 - At most $3L$ requests can be served.
 - The expected number of requests ALG missed is at least $F/4$.
 - OPT' served $3L$ type A requests.

Therefore, $\text{OPT} \geq \text{OPT}' = 3L$ and the competitive ratio depends only on F and L :

$$\begin{aligned} \frac{\text{OPT}(\sigma)}{E(\text{ALG}(\sigma))} &\geq \frac{3L}{3L - \frac{1}{4}F} = \frac{3L}{3L - \frac{1}{4}(\sqrt{w(S)L})} \\ &= \frac{3L}{3L - \frac{1}{4}(\sqrt{MST(G)L})} = 1 + \Omega\left(\sqrt{TSP(G)/L}\right). \end{aligned}$$

Here the second equality results from the fact that $w(S) = MST(G)$.

3. In Termination Case 3 ALG served type A request and at least $F/2$ type B requests at each block. In Theorem 4.2 we proved that:

- At most $3L$ requests can be served.
- The expected number of requests ALG missed in each block due to traveling time is at least $\frac{1}{8} \frac{w(S)}{2} = \frac{w(S)}{16}$.
- OPT' served $3L$ type A requests. Therefore $\text{OPT} \geq \text{OPT}' = 3L$.

By the first property required by this theorem, each sequence of nodes in G requires more traveling time than in S . Therefore, the expected number of requests ALG missed per block is at least $\frac{w(S)}{16}$. Since the number of blocks is $\frac{1}{3} \sqrt{\frac{L}{w(S)}}$, we conclude that the competitive ratio is:

$$\begin{aligned} \frac{\text{OPT}(\sigma)}{E(\text{ALG}(\sigma))} &\geq \frac{3L}{3L - \left(\frac{1}{3} \sqrt{L/w(S)}\right) \frac{w(S)}{16}} \\ &= \frac{3L}{3L - \left(\frac{1}{3} \sqrt{L/\text{MST}(G)}\right) \frac{\text{MST}(G)}{16}} = 1 + \Omega\left(\sqrt{\text{TSP}(G)/L}\right). \end{aligned}$$

Here the first equality result from the fact that $w(S) = \text{MST}(G)$.

Note that in all 3 cases we got $1 + \Omega\left(\sqrt{\text{TSP}(G)/L}\right) = 1 + \Omega(\sqrt{\delta})$.

This completes the proof of the Theorem 4.4.

A.2 Proof of Theorem 5.1

First we need to demonstrate that the output ρ is feasible. Specifically, we need to prove that every request i is served during the time window $[r_i, d_i]$, and that there is a traveling time of length $w(v_i, v_j)$ between serving any two successive requests with different nodes v_i and v_j .

Lemma A.1 *The algorithm ORIENT-WINDOW generates a valid output.*

Proof. All request in R^ℓ have deadlines of at least $K\ell$. Hence, when they are served they have not been expired yet. Serving at phase ℓ ends at time $K\ell$ even if not all requests in R^ℓ are served. Hence, the phases of the algorithm are well defined. \square

Now we analyze the performance guarantee of the algorithm. Let σ' be the following modification of σ . It consists of all requests in σ such that a request $([r, d], v) \in \sigma$ is replaced by a request $([K\lceil r/K \rceil, K\lfloor d/K \rfloor], v) \in \sigma'$. Hence, the release and deadline times of each request in σ' are aligned with the start/end time of the corresponding phase so that the time window of each request is fully contained in the time window of that request according to σ .

The notion of λ -perturbation, defined in [3], is as follows: An input sequence $\hat{\delta}$ is a λ -perturbation of δ if $\hat{\delta}$ consists of all requests of δ , and each request $[\hat{r}, \hat{d}] \in \hat{\delta}$ corresponding to request $[r, d] \in \delta$ satisfies $\hat{r} - r \leq \lambda$ and $d - \hat{d} \leq \lambda$.

By definition, σ' is K -perturbation of σ .

We use Theorem 2.2 from [3]:

Theorem A.2 *Suppose $\hat{\delta}$ is a λ -perturbation of δ then $\text{OPT}(\hat{\delta}) = (1 - 2\lambda/L)\text{OPT}(\delta)$ (where L is the minimum laxity).*

By applying the Theorem we get

$$\text{OPT}(\sigma') \geq \left(1 - 2\frac{3\Delta(G)}{9\Delta(G)}\right)\text{OPT}(\sigma) = \text{OPT}(\sigma)/3. \quad (1)$$

Let OPT' be an optimal offline algorithm that is not allowed to serve any request that was served by the online algorithm. It is clear that $\text{OPT}(\sigma) \leq \text{OPT}'(\sigma) + \text{ALG}(\sigma)$. Moreover, at each time unit the set of the unserved requests of $\text{OPT}'(\sigma)$ is a subset of the unserved requests of $\text{ALG}(\sigma)$.

Lemma A.3 $\text{ALG}(\sigma') \geq \text{OPT}'(\sigma')/9$

Proof. We prove the Lemma for each phase separately. Recall that in each specific phase ℓ we compute 3-approximation to the unrooted orienteering problem [4] with budget $\Delta(G)$ where the prize of a node v_j is the number of requests in S_j^ℓ (actually there is a $(2 + \epsilon)$ -approximation [8] but for simplicity we use the 3-approximation). Let x be the total prize in the solution for phase ℓ . We separate into two cases. If $x \geq \Delta(G)$ then ALG serves at least $\Delta(G)$ requests. It is clear that he needs at most $3\Delta(G)$ time units to serve the requests (at most $\Delta(G)$ time units to travel to the first node in the solution, at most $\Delta(G)$ time units to travel between the nodes of the solution and $\Delta(G)$ time units to serve the requests). Since the size of each phase is $3\Delta(G)$, there are enough time units to serve $\Delta(G)$ requests. OPT' can serve up to $3\Delta(G)$ requests (a request per time unit). Hence, $\text{ALG}(\sigma') \geq \text{OPT}'(\sigma')/3$ in phase ℓ .

If $x < \Delta(G)$ ALG serves x requests follows by a similar argument to the one described in the previous case. Assume by a contradiction that OPT' served more than $9x$ requests. Since the size of the phase is $3\Delta(G)$ it is clear that there exists $\Delta(G)$ consecutive time units during the phase in which OPT' served more than $9x/3 = 3x$ requests. This is a contradiction to the correctness of the approximation algorithm (recall that OPT' is not allowed to serve requests that were served by ALG . Therefore, the set of the unserved requests of $\text{OPT}'(\sigma)$ is a subset of the unserved requests of $\text{ALG}(\sigma)$). Hence, OPT' served up to $9x$ requests and $\text{ALG}(\sigma') \geq \text{OPT}'(\sigma')/9$ in phase ℓ . \square

Now are ready to prove that ORIENT-WINDOW attains a competitive ratio of $O(1)$.

Using the previously stated results, we obtain that

$$\text{ALG}(\sigma) = \text{ALG}(\sigma') \geq \text{OPT}'(\sigma')/9 \geq \text{OPT}'(\sigma)/(9 * 3) = \text{OPT}'(\sigma)/27$$

The first equality follows by the definition of the algorithm. The first inequality results from Lemma A.3. The second inequality holds by Equation (1).

Combining the result with the relation between $\text{OPT}'(\sigma)$ and $\text{OPT}(\sigma)$ we get:

$$\text{OPT}(\sigma) \leq \text{OPT}'(\sigma) + \text{ALG}(\sigma) \leq 28\text{ALG}(\sigma)$$

This complete the proof.

A.3 Proof of Theorem 5.2

First we need to demonstrate that the output ρ is feasible. Specifically, we need to prove that every request i is served during the time window $[r_i, d_i]$, and that there is a traveling time of length $w(v_i, v_j)$ between serving any two successive requests with different nodes v_i and v_j .

Lemma A.4 *The algorithm TSP-EDF generates a valid output.*

Proof. The proof is similar to the proof of Lemma A.1. \square

Now we analyze the performance guarantee of the algorithm. First we define two input sequences σ' and $\tilde{\sigma}$, which are modifications of σ . The input sequence σ' consists of all requests in σ , but modifies the node of the requests to a fixed node v' . Specifically, each request $([r, d], v) \in \sigma$ defines a request $([r, d], v') \in \sigma'$. The input sequence $\tilde{\sigma}$ consists of all requests in σ such that a request $([r, d], v) \in \sigma$ is replaced by a request $([K\lceil r/K \rceil, K\lfloor d/K \rfloor], v') \in \tilde{\sigma}$, where v' is a fixed node. Hence, all requests in $\tilde{\sigma}$ have the same node, and the release and deadline times of each request in $\tilde{\sigma}$ are aligned with the start/end time of the corresponding phase so that the time window of each request is fully contained in the time window of that request according to σ .

Lemma A.5 $\text{OPT}(\tilde{\sigma}) = \text{ALG}(\tilde{\sigma})$.

Proof. Note that algorithm TSP-EDF has three modification with respect to EDF:

- request deadline times are modified to $K\lfloor d/K \rfloor$.
- request release times are modified to $K\lceil r/K \rceil$ (because in each phase only requests released during previous phases are served).
- Traveling time units are added between serving requests from different nodes.

The release and deadline times of the requests in $\tilde{\sigma}$ are aligned and all the requests have the same node. Hence, ALG's schedule is identical to EDF's schedule. Since EDF is optimal for sequences that consist of requests with one node (observe that this is the same as scheduling packets with unit size and unit value), $\text{OPT}(\tilde{\sigma}) = \text{ALG}(\tilde{\sigma})$. \square

By definition, $\tilde{\sigma}$ is K -perturbation of σ' , and the nodes of all requests are identical. Hence, Theorem 2.2 from [3] (see Section A.2 for the definitions of perturbation and Theorem 2.2) yields the following inequality:

$$\text{OPT}(\tilde{\sigma}) \geq \left(1 - 2\sqrt{\text{TSP}(\text{G})/L}\right) \text{OPT}(\sigma'). \quad (2)$$

Lemma A.6 $\text{ALG}(\sigma) \geq \left(1 - \sqrt{\text{TSP}(\text{G})/L}\right) \text{ALG}(\sigma')$.

Proof. The difference between the schedule TSP-EDF generated for σ and the schedule it generates for σ' is that requests might be dropped at the end of each phase in σ due to traveling time. The worst case for σ is when there are no idle time units in any of the phases of σ' . Otherwise, the idle time units might be used for traveling time. Therefore, there are at most $\lceil \text{ALG}(\sigma')/K \rceil - 1$ phases in which algorithm TSP-EDF serves requests (the -1 term is due to the fact that the algorithm does not serve any request during the first phase). Since there are at most $\text{TSP}(\text{G})$ traveling time units in each phase, we obtain the following inequality:

$$\begin{aligned} \text{ALG}(\sigma) &\geq \text{ALG}(\sigma') - (\lceil \text{ALG}(\sigma')/K \rceil - 1) \text{TSP}(\text{G}) \\ &= \text{ALG}(\sigma') - \left(\left\lceil \frac{\text{ALG}(\sigma')}{\sqrt{\text{TSP}(\text{G})L}} \right\rceil - 1 \right) \text{TSP}(\text{G}) \\ &\geq \left(1 - \sqrt{\text{TSP}(\text{G})/L}\right) \text{ALG}(\sigma'). \end{aligned}$$

\square

Now are ready to prove that TSP-EDF attains a competitive ratio of $1 + O(\sqrt{\text{TSP}(\text{G})/L})$.

Using the previously stated results, we obtain that

$$\begin{aligned}
\text{ALG}(\sigma) &\geq \left(1 - \sqrt{\text{TSP}(\text{G})/L}\right) \text{ALG}(\sigma') = \left(1 - \sqrt{\text{TSP}(\text{G})/L}\right) \text{ALG}(\tilde{\sigma}) \\
&= \left(1 - \sqrt{\text{TSP}(\text{G})/L}\right) \text{OPT}(\tilde{\sigma}) \\
&\geq \left(1 - \sqrt{\text{TSP}(\text{G})/L}\right) \left(1 - 2\sqrt{\text{TSP}(\text{G})/L}\right) \text{OPT}(\sigma') \\
&\geq \left(1 - 3\sqrt{\text{TSP}(\text{G})/L}\right) \text{OPT}(\sigma).
\end{aligned}$$

The first inequality results from Lemma A.6. The first equality follows by the definition of the algorithm. The second equality holds by lemma A.5. The second inequality results from Equation (2). Finally, the last inequality holds because σ' is similar σ , but all requests have the same node. This implies that any schedule feasible for σ is also feasible for σ' , and thus $\text{OPT}(\sigma') \geq \text{OPT}(\sigma)$.

We conclude that $\frac{\text{OPT}(\sigma)}{\text{ALG}(\sigma)} \geq 1 + O(\sqrt{\text{TSP}(\text{G})/L})$

B Figures




Figure 4: Block's structure. The pair (r, d) represent release time r and deadline d . Note that all the requests arrive at once in the beginning of the block.

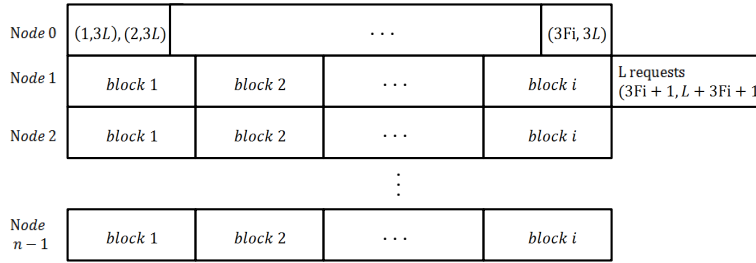


Figure 5: Sequence structure for Termination Case 1. See Figure 4 for blocks structure.

Node 0	(1,3L), (2,3L)	...			(3Fi, 3L)	3L requests (3Fi + 1,3L) Final event
Node 1	block 1	block 2	...	block i		
Node 2	block 1	block 2	...	block i		
			⋮			
Node $n - 1$	block 1	block 2	...	block i		

Figure 6: Sequence structure for Termination Case 2. See Figure 4 for blocks structure.

Node 0	(1,3L),(2,3L)	...			(L,3L)	3L requests (L + 1,3L)
Node 1	block 1	block 2	...	block N	Final event	
Node 2	block 1	block 2	...	block N		
⋮						
Node n - 1	block 1	block 2	...	block N		

Figure 7: Sequence structure for Termination Case 3. Recall that $N = \frac{1}{3}\sqrt{\frac{L}{w(S)}}$. See Figure 4 for blocks structure.